

Drawing Partially Embedded and Simultaneously Planar Graphs^{*}

Timothy M. Chan¹, Fabrizio Frati², Carsten Gutwenger³, Anna Lubiw¹,
Petra Mutzel³, and Marcus Schaefer⁴

¹ Cheriton School of Computer Science, University of Waterloo, Canada.
{tmchan, alubiw}@uwaterloo.ca

² School of Information Technologies, The University of Sydney, Australia.
fabrizio.frati@sydney.edu.au

³ Technische Universität Dortmund, Dortmund, Germany.
{carsten.gutwenger, petra.mutzel}@tu-dortmund.de

⁴ DePaul University, Chicago, Illinois, USA.
mschaefer@cdm.depaul.edu

Abstract. We investigate the problem of constructing planar drawings with few bends for two related problems, the *partially embedded graph* problem—to extend a straight-line planar drawing of a subgraph to a planar drawing of the whole graph—and the *simultaneous planarity* problem—to find planar drawings of two graphs that coincide on shared vertices and edges. In both cases we show that if the required planar drawings exist, then there are planar drawings with a linear number of bends per edge and, in the case of simultaneous planarity, a constant number of crossings between every pair of edges. Our proofs provide efficient algorithms if the combinatorial embedding of the drawing is given. Our result on partially embedded graph drawing generalizes a classic result by Pach and Wenger which shows that any planar graph can be drawn with a linear number of bends per edge if the location of each vertex is fixed.

1 Introduction

In many practical applications we wish to draw a planar graph while satisfying some geometric or topological constraints. One natural situation is that we have a drawing of part of the graph and wish to extend it to a planar drawing of the whole graph. Pach and Wenger [25] considered a special case of this problem. They showed that any planar graph can be drawn with its vertices lying at pre-assigned points in the plane and with a linear number of bends per edge. In this case the pre-drawn subgraph has no edges.

If the pre-drawn subgraph H has edges, a planar drawing of the whole graph G extending the given drawing \mathcal{H} of H may not exist. Angelini et al. [1] gave a linear-time algorithm for the corresponding decision problem; the algorithm returns, for a positive answer, a planar embedding of G that *extends* that of \mathcal{H}

^{*} A preliminary version of this paper appeared in [7].

(i.e., if we restrict the embedding of G to the edges and vertices of H , we obtain the embedding corresponding to \mathcal{H}). If one does not care about maintaining the actual planar drawing of H this is the end of the story, since standard methods can be used to find a straight-line planar drawing of G in which the drawing of H is topologically equivalent to the one of \mathcal{H} . In this paper we show how to draw G while preserving the actual drawing \mathcal{H} of H , so that each edge has a linear number of bends. This bound is worst-case optimal, as proved by Pach and Wenger [25] in the special case in which H has no edges.

A result analogous to ours was claimed by Fowler et al. [14] for the special case in which H has the same vertex set as G . Their algorithm draws the edges of G one by one, in any order so that edges connecting distinct connected components of H precede edges within the same connected component of H ; each edge is drawn as a curve with the minimum number of bends. Fowler et al. claim that their algorithm constructs drawings with a linear number of bends per edge. However, we prove that there exists a tree, a planar drawing of its vertex set, and an order of the edges of the tree, such that drawing the edges in the given order as curves with the minimum number of bends results in some edges having an exponential number of bends.

The second graph drawing problem we consider is the *simultaneous planarity* problem [5], also known as “simultaneous embedding with fixed edges (SEFE)”. The SEFE problem is strongly related to the partially embedded graph problem and—in a sense we will make precise later—generalizes it. We are given two planar graphs G_1 and G_2 that share a *common subgraph* G (i.e., G is composed of those vertices and edges that belong to both G_1 and G_2). We wish to find a *simultaneously planar drawing*, i.e., a planar drawing of G_1 and a planar drawing of G_2 that coincide on G . Graphs G_1 and G_2 are *simultaneously planar* if they admit such a drawing. Both G_1 and G_2 may have *private* edges that are not part of G . In a simultaneous planar drawing the private edges of G_1 may cross the private edges of G_2 . The simultaneous planarity problem arises in information visualization when we wish to display two relationships on two overlapping element sets.

The decision version of the simultaneous planarity problem is not known to be **NP**-complete, or to be solvable in polynomial time, though it is known to be **NP**-complete if more than two graphs are given [15]. However, there is a combinatorial characterization of simultaneous planarity, based on the concept of a “compatible embedding”, due to Jünger and Schulz [20] (see below for details). Erten and Kobourov [12], who first introduced the problem, gave an efficient drawing algorithm for the special case where the two graphs share vertices but no edges. In this case, a simultaneous planar drawing on a polynomial-size grid always exists in which each edge has at most two bends and therefore any two edges cross at most nine times, see [11,12,21]. In this paper we show that if two graphs have a simultaneous planar drawing, then there is a drawing on a polynomial-size grid in which every edge has a linear number of bends and in which any two edges cross at most 24 times. Our result is algorithmic, assuming a compatible embedding is given.

1.1 Realizability Results

Our paper addresses the following two drawing problems:

Planarity of a partially embedded graph (PEG). Given a planar graph G and a straight-line planar drawing \mathcal{H} of a subgraph H of G , find a planar drawing of G that extends \mathcal{H} (see [1,19]).

Simultaneous planarity (SEFE). Given two planar graphs G_1 and G_2 that share a subgraph G , find a simultaneous planar drawing of G_1 and G_2 (see [5]).

We prove the following results:

Theorem 1 (Realizing a Partially Embedded Graph). *Let G be an n -vertex planar graph, let H be a subgraph of G , and let \mathcal{H} be a straight-line planar drawing of H . Suppose that G has a planar embedding \mathcal{E} that extends the one of \mathcal{H} . Then we can construct a planar drawing of G in $O(n^2)$ -time which realizes \mathcal{E} , extends \mathcal{H} , and has at most $72|V(H)|$ bends per edge.*

Theorem 1 generalizes Pach and Wenger’s classic result, which corresponds to the special case in which the pre-drawn subgraph has no edges.

Theorem 2 (Realizing a Simultaneous Planar Embedding). *Let G_1 and G_2 be simultaneously planar graphs on a total of n vertices with a shared subgraph G . If we are given a compatible embedding of the two graphs, then we can construct in $O(n^2)$ time a drawing that realizes the compatible embedding, and in which any private edge of G_1 and any private edge of G_2 intersect at most 24 times. In addition, we can ensure either one of the following two properties:*

- (i) *each edge of G is straight, and each private edge of G_1 and of G_2 has at most $72n$ bends; also, vertices, bends, and crossings lie on an $O(n^2) \times O(n^2)$ grid; or*
- (ii) *each edge of G_1 is straight and each private edge of G_2 has at most $72|V(G_1)|$ bends per edge.*

Theorem 1 provides a weak form of Theorem 2: If G_1 and G_2 are simultaneously planar, they admit a compatible embedding. Take any straight-line planar drawing of G_1 realizing that embedding and extend the induced drawing of G to a drawing of G_2 . By Theorem 1, we obtain a simultaneous planar drawing where each edge of G_1 is straight and each private edge of G_2 has at most $72|V(G_1)|$ bends per edge. Our stronger result of 24 crossings between any two edges is obtained by modifying the proof of Theorem 1, rather than applying that result directly.

Grilli et al. [16] very recently and independently proved a result in some respect stronger than Theorem 2. They showed that two simultaneously planar graphs have a simultaneous planar drawing with at most 9 bends per edge, vastly better than our $72n$ bound. On the other hand, our bound of 24 crossings per pair of edges is better than the bound of 100 that can be derived from their

result. Also, our algorithm allows us to construct simultaneous planar drawings in which each edge of one graph is straight or in which vertices, bends, and crossings lie on a polynomial-size grid. The former feature is not achievable by means of Grilli et al.’s algorithm; the latter one could be obtained from Grilli et al.’s result, at the expense of increasing the number of bends per edge to $300n$ (which corresponds to the number of crossings on a single private edge).

1.2 Related Work

The decision version of simultaneous planarity generalizes partially embedded planarity: given an instance (G, H, \mathcal{H}) of the latter problem, we can augment \mathcal{H} to a drawing of a 3-connected graph G_1 and let $G_2 = G$. Then G_1 and G_2 are simultaneously planar if and only if G has a planar embedding extending \mathcal{H} . In the other direction, the algorithm [1] for testing planarity of partially embedded graphs solves the special case of the simultaneous planarity problem in which the embedding of the common graph G is fixed (which happens, e.g., if G or one of the two graphs is 3-connected).

Several optimization versions of partially embedded planarity and simultaneous planarity are **NP**-hard. Patrignani showed that testing whether there is a straight-line drawing of a planar graph G extending a given drawing of a subgraph of G is **NP**-complete [26], so bend minimization in partial embedding extensions is **NP**-complete; Patrignani’s result holds even if a combinatorial embedding of G is given.⁵ Bend minimization in simultaneous planar drawings is **NP**-hard, since it is **NP**-hard to decide whether there is a straight-line simultaneous drawing [13]. Crossing minimization in simultaneous planar drawings is also **NP**-hard, as follows from an **NP**-hardness result on *anchored planar drawings* by Cabello and Mohar [6]; see Theorem 4 in Section 4 for a slightly stronger result.

Di Giacomo et al. [10] studied the special case of PEG in which the n -vertex graph G to be drawn is a tree. They showed that, given a drawing \mathcal{H} of a subtree H of G , a drawing of G extending \mathcal{H} can be computed in $O(n^2 \log n)$ time so that each edge of G has at most $1 + 2\lceil |V(H)|/2 \rceil$ bends.

Further, as mentioned above, the special cases of PEG and SEFE in which there are no edges in the pre-drawn subgraph and in the common subgraph have been already studied.

Concerning PEG, Pach and Wenger [25] proved the following result: given an n -vertex planar graph G with fixed vertex locations, a planar drawing of G in which each edge has at most $120n$ bends can be constructed in $O(n^2)$ time. They also proved that such a bound is asymptotically tight in the worst case. Regarding the constant, Badent et al. [2] improved the bound to $3n+2$ bends per edge. Biedl and Floderus [4] considered the more general problem of drawing an n -vertex planar graph on fixed vertex locations where the drawing is constrained

⁵ Patrignani does not explicitly claim **NP**-completeness in the case in which the embedding of G is fixed, but that can be concluded by checking his construction; only the variable gadget, pictured in his Figure 3, needs minor adjustments.

to lie inside a k -vertex polygon. They show that there is a drawing with $O(n+k)$ bends per edge.

Concerning SEFE, Di Giacomo and Liotta [11] and independently Kammer [21] proved the following result: given two planar graphs G_1 and G_2 sharing some vertices and no edge with a total number of n vertices, there exists an $O(n)$ -time algorithm to construct a simultaneous planar drawing of G_1 and G_2 on a grid of size $O(n^2) \times O(n^2)$, where each edge has at most 2 bends, hence there are at most 9 crossings between any edge of G_1 and any edge of G_2 . This improves upon a previous result of Erten and Kobourov [12]. The algorithms in [11,12,21] make use of a drawing technique introduced by Kaufmann and Wiese [22].

Haeupler et al. [17] showed that if two simultaneously planar graphs G_1 and G_2 share a subgraph G that is connected, then there is a simultaneous planar drawing in which any edge of G_1 and any edge of G_2 intersect at most once. Introducing vertices at crossing points yields a planar graph, and a straight-line drawing of that graph provides a simultaneous planar drawing with $O(n)$ bends per edge, $O(n)$ crossings per edge, and with vertices, bends, and crossings on an $O(n^2) \times O(n^2)$ grid. Our result generalizes this to the case where the common graph G is not necessarily connected.

1.3 Graph Drawing Terminology

A *rotation system* for a graph is a cyclic ordering of the edges incident to each vertex. A rotation system of a connected graph determines its *facial walks*—the closed walks in which each edge (u, v) is followed by the next edge (v, w) in the cyclic order at v . The *size* $|W|$ of a facial walk W is the number of vertices on W , where we count vertex repetitions. (Note that a graph that consists of a single vertex has a single facial walk of size 1; for any other connected graph the size of a facial walk is equal to the number of edges in the facial walk, counting repetitions.) A rotation system is *planar* if Euler’s formula holds, i.e., $n - m + f = 2$ where n is the number of vertices, m is the number of edges, and f is the number of facial walks. A *planar embedding* of a graph consists of a planar rotation system together with a specified outer face. A fundamental result about connected planar graphs is that every planar drawing corresponds to a planar embedding, and conversely, every planar embedding can be realized as a planar drawing (and, in fact, as a straight-line planar drawing by Fáry’s theorem). Furthermore, facial walks correspond to faces in the drawing.

These definitions do not handle the combinatorics of a planar drawing of a disconnected graph—namely the definition of planar embedding as stated above does not tell us how connected components nest into each other.

Following Jünger and Schulz [20], we define a *topological embedding* of a (possibly non-connected) graph as follows: We specify a planar embedding for each connected component. This determines a set of inner faces. For each connected component we specify a “containing” face, which may be an inner face of some other component or the unique outer face. Furthermore, we forbid cycles of containment—in other words, if a connected component is contained in an inner

face, which is contained in a component, etc., then this chain of containments must lead eventually to the unique outer face.

A *facial boundary* in a topological embedding of a graph is the collection of facial walks along the (not necessarily connected) boundary of a face. Each face (unless it is the outer face) has a distinguished facial walk we call the *outer* facial walk separating the remaining *inner* facial walks from the outer face of the embedding. The *size* of a facial boundary is the sum of the sizes of its facial walks.

A *compatible embedding* of two planar graphs G_1 and G_2 consists of topological embeddings of G_1 and G_2 such that the common subgraph G inherits the same topological embedding from G_1 as from G_2 (where a subgraph inherits a topological embedding in a straightforward way; in particular, if we remove an edge that disconnects the graph, the face containment is determined by the edge that was removed). Jünger and Schulz [20] proved that G_1 and G_2 are simultaneously planar if and only if they have a compatible embedding. For that proof, they construct a simultaneous planar drawing of G_1 and G_2 by extending a drawing of G (thus proving a form of our Theorem 1). However, their method does not yield any bounds on the number of bends or crossings.

2 Partially Embedded Graphs

In this section we prove Theorem 1; that is, we show how to construct a planar drawing of G that extends the planar straight-line drawing \mathcal{H} and has a linear number of bends per edge assuming that we are given a planar embedding of G extending \mathcal{H} . It is sufficient to prove the result for a single face F of \mathcal{H} , since the embedding of G is given, and we know for each vertex and edge of G which face of \mathcal{H} it lies in, so the drawings in different faces of \mathcal{H} do not interfere with each other.

Pach and Wenger [25] proved their upper bound on the number of bends needed to draw a graph with fixed vertex locations by drawing a tree with its leaves at the fixed vertex locations, and “routing” all the edges close to the tree, sometimes crossing the tree but never crossing each other. We want to use their approach, but we have to deal with a more general problem. Instead of fixed vertex locations we have fixed facial boundaries. The solution is natural: We contract each facial walk W_i of F to a single vertex v_i , fix a position for vertex v_i inside F near W_i , and then apply the Pach-Wenger method to draw the contracted graph on the fixed vertex locations v_i . We ensure that the contracted graph is drawn inside F , indeed we stay a small distance away from the boundary of F , inside a polygonal region F' that is an “inner approximation” of F . Inside F' we draw a tree T with its leaves v_i at the fixed vertex locations, while suitably bounding the size of T so as to get our bound on the number of bends. We then route the edges of the contracted graph close to T as Pach and Wenger do. Finally, to retrieve the original, uncontracted graph, we route the edges incident to v_i to their true endpoint on the facial boundary W_i —these routes use the empty buffer zone between F and F' .

We fill in the details of this argument in Section 2.3, but before doing so we introduce “inner approximations” in Section 2.1, and formalize the tree argument in Section 2.2.

To simplify notation, we use n_A and m_A for the number of vertices and edges in a graph (or subgraph) A .

2.1 Approximating Faces

In the drawing \mathcal{H} , the face F is a region of the plane homeomorphic to a disc with holes. Each facial walk of F appears in the drawing as a *closed polygonal arc*, i.e. a sequence of straight-line segments joined in a path that returns to its starting point (repeated segments/vertices may occur). We will refer to a facial walk and its drawing interchangeably.

We will approximate F by offsetting each of its facial walks into the interior of F . See Figure 1. Let W_1 be the outer facial walk of F , and let W_2, \dots, W_b be the inner facial walks. An *inner ε -approximation* of W_i is a simple polygon P_i (a closed polygonal arc with no self-intersections) such that:

1. P_i is ε -close to W_i , meaning that every point of P_i is within distance ε of a point of W_i ,
2. the inner facial walk W_i lies in the interior of P_i , for each $2 \leq i \leq b$, and
3. the outer facial walk W_1 lies in the exterior of P_1 .

If in addition the P_i ’s form a *polygonal region* (a simple polygon with holes) with P_1 as the outer polygon, then we say that the polygonal region is an *inner ε -approximation* of F . The next lemma shows that we can build inner ε -approximations of F .

Lemma 1. *For any $\varepsilon > 0$ we can efficiently construct an inner ε -approximation F' of F .*

See Figure 1 for an illustration of Lemma 1. To prove the lemma, we construct—for every sufficiently small $\varepsilon > 0$ and for every facial walk of F —an inner ε -approximating polygon P_ε which does not have too many bends, and so that the P_ε are *nested* in the following sense: if $0 < \varepsilon' < \varepsilon$, then $P_{\varepsilon'}$ lies in the interior of P_ε if F is an inner face, and vice versa otherwise. There are various ways to achieve this. Pach and Wenger [25] use the Minkowski sum of the facial walk (in their case the facial walk of a tree) and a square diamond centered at 0. We use a slightly different construction, because it seems easier (both computationally and conceptually) and it gives a slightly better bound on the number of bends (which is what we are most interested in): for the facial walk of an n -vertex tree, Pach and Wenger construct a polygon with $4n - 2$ vertices, while ours have $2n - 2$ vertices. Our construction does have one disadvantage: the resulting drawings are tight, placing elements close together, for sharp (acute or obtuse) angles (the Minkowski-sum construction has the same problem for highly obtuse angles only).




Fig. 1. (a) A face F with outer facial walk W_1 and inner facial walks W_2, W_3, W_4 . (b) An inner approximation F' (heavy blue lines).

Lemma 2. *Let W be a facial walk in a face F of a drawing of a graph G in the plane. We can efficiently construct a nested family of inner ε -approximating polygons P_ε so that each P_ε has at most $\max\{3, |W|\}$ vertices.*

Proof. Let e, v, f be a corner of W , that is, two consecutive edges e, f and their shared vertex v . At v erect the angle bisector of e and f of length ε (inside F), and let v' be the endpoint of the bisector different from v . For computational reasons, it may be better to use the ℓ_1 -norm at this point (the Euclidean norm will lead to square root expressions in the coordinates). If $(v_i)_{i=1}^k$ is the sequence of vertices along W , with $k = |W|$, then $(v'_i)_{i=1}^k$ defines a closed polygonal chain. If ε is sufficiently small, namely less than half the distance between any vertex of W and a non-adjacent edge on W , the polygonal chain is free of self-crossings, and therefore bounds a simple polygon with $|W|$ vertices. There are two special cases in which this argument does not work: if the facial walk is a facial walk on an isolated vertex or an isolated edge. In both of these cases, we can approximate W using a triangle.

To prove Lemma 1 we can use Lemma 2 to efficiently construct an inner ε -approximating polygon for each facial walk of F . The resulting polygons are disjoint and form a polygonal region as long as ε is less than half the distance between any two non-adjacent vertices or edges of \mathcal{H} .

2.2 Extending Partial Embeddings

Our main technical tool in the proof of Theorem 1 is the following lemma. We suggest skipping the proof of this lemma in a first reading. Multigraphs, in this paper, may have multiple edges and loops.

Lemma 3. *Let G be a multigraph with a given planar embedding and fixed locations for a subset U of its vertices. Suppose we are given a straight-line drawing of a tree T whose leaves include all the vertices in U at their fixed locations.*

Then for every $\varepsilon > 0$ there is a planar poly-line drawing of G that is ε -close to T , that realizes the given embedding, where the vertices in U are at their fixed locations, where each edge has at most $12n_T$ bends, and where each edge comes close to each vertex u in U at most six times (where coming close to u means entering and leaving an ε -neighborhood of u or terminating at u).

Our proof of Lemma 3 will follow closely the structure of Pach and Wenger's algorithm [25] to draw a planar graph with fixed vertex locations. That algorithm has three ingredients: (i) making G Hamiltonian, (ii) drawing the Hamiltonian cycle of G , and (iii) drawing the remaining edges of G . We use their result (i) directly:

Lemma 4 (Pach, Wenger [25]). *Given a planar graph G we can in linear time construct a Hamiltonian graph G' with $|E(G')| \leq 5|E(G)| - 10$ by adding and subdividing edges of G (each edge is subdivided by at most two new vertices).*

We will use a slightly stronger version of Lemma 4 in which G is allowed to be a multigraph. Pach and Wenger's proof of Lemma 4 works in the presence of multiple edges and loops.

For part (ii) Pach and Wenger show that a Hamiltonian cycle can be drawn at fixed vertex locations ε -close to a star connecting all the vertices. For our application, we replace their star with a straight-line drawing of a tree T whose leaves are the vertices v_i . Lemma 5 shows how to draw the Hamiltonian cycle. Later we will see how to draw the remaining edges.

Independently of our result, the generalization of part (ii) to trees has essentially been shown by Chan et al. [8]. Since their goal was to minimize edge lengths, they did not give an estimate on the number of bends.

Lemma 5. *Let C be a cycle with fixed vertex locations, and suppose we are given a straight-line planar drawing of a tree T , in which the vertices of C are leaves of T at their fixed locations. Then for every $\varepsilon > 0$ there is a planar poly-line drawing of C with at most $2|E(T)| - 1$ bends per edge and ε -close to T .*

Proof. Let p_1, \dots, p_n be the vertices of C in their order along the cycle. We build a planar poly-line drawing of C as follows. Let Θ_i be an $i\varepsilon/n$ -approximation of T for $1 \leq i < n$ (which we construct using Lemma 2). We start at p_1 . Suppose we have already built the poly-line drawing of p_1, \dots, p_i and we want to add $p_i p_{i+1}$. Let Q_i be the unique path in T connecting p_i to p_{i+1} . Create Θ'_i from Θ_i by keeping only the vertices of Θ_i close to (approximating) vertices in $T_i := \bigcup_{j \leq i} Q_j$. This removes parts of the walk along Θ_i which we patch up as follows: suppose v is an interior vertex of T_i , and v is incident to e which does not lie on T_i . Then v is approximated by two vertices v_1 and v_2 which lie on bisectors formed by e with neighboring edges. Now v_1 and v_2 belong to Θ'_i , but the path along Θ_i between them got removed (since e does not belong to T_i). We add $v_1 v_2$ to Θ'_i to connect them. Note that $v_1 v_2$ does not pass through v since v is incident to at least three edges (e and two edges of T_i), and it does not cross any edges of any Θ'_j with $j < i$, since T_i is monotone: if $e \notin E(\Theta_i)$, then $e \notin E(\Theta_j)$ for $j < i$.




Fig. 2. The underlying tree T is in black (thick edges), angle bisectors in gray; the Θ'_i are drawn as thin black edges; to reduce clutter, we are not showing the remaining edges of Θ_i ; the drawing of C is indicated by the green line.

See Figure 2 for an illustration. Now both p_i and p_{i+1} correspond to unique vertices on Θ'_i (since they are leaves), so we can pick the facial walk v_1, \dots, v_k on Θ'_i which connects p_i to p_{i+1} and which avoids passing by p_1 . We now add line segments $p_i v_2, v_2 v_3, \dots, v_{k-2} v_{k-1}, v_{k-1} p_{i+1}$ to the poly-line drawing of C . We treat the final edge $p_n p_1$ similarly, except that we move along Θ'_{n-1} back to p_1 in the last step, which we can do, since none of the intermediate paths passed by p_1 . Each edge of C is replaced by a polygonal arc with at most $2|E(T)| - 1$ bends.

The following lemma shows how to draw the remaining edges of G , assuming that G is Hamiltonian. As mentioned earlier, this lemma is close to a result by Chan et al. [8], except for the claim about the number of bends, and the rotation system (which we need for our main result).

Lemma 6. *Let G be a Hamiltonian multigraph with a given planar embedding and fixed vertex locations. Suppose we are given a straight-line drawing of a tree T whose leaves include all the vertices of G at their fixed locations. Then for every $\varepsilon > 0$ there is a planar poly-line drawing of G that is ε -close to T , that realizes the given embedding, and so that the vertices of G are at their fixed locations, every edge has at most $4|E(T)| - 1$ bends, and every edge comes close to any leaf of T at most twice.*

The obvious idea—routing edges along the Hamiltonian cycle C —only gives a quadratic bound on the number of bends, since each edge would follow the path of a linear number of edges of C , and each edge of C has a linear number of

bends. Pach and Wenger came up with an ingenious way to construct auxiliary curves with few bends based on the level curves Θ'_i which carry the cycle C in the proof of Lemma 5.

Proof. Let C be the Hamiltonian cycle of G and let G_1 and G_2 be the two outerplanar graphs composed of C and, respectively, of the edges of G outside and inside C . Using Lemma 5 we find a planar poly-line drawing of C on $V(G)$. We need to show how to draw G_1 and G_2 respecting the planar embeddings induced by the given embedding of G . Let $n = |V(G)|$ and $m_i = |E(G_i)|$. We only describe how to draw G_1 , since G_2 can be handled analogously. Let $\Delta_{i,k}$, $1 \leq k \leq m_1$ be a $k\varepsilon/(nm_1)$ -approximation of Θ'_i constructed using Lemma 2. For a fixed i , each $\Delta_{i,k}$ crosses C twice: when C moves from p_i to Θ'_{i+1} , and when it finally moves back from Θ'_n to p_1 . As in Pach and Wenger, we can then split $\Delta_{i,k}$ at the crossings and connect their free ends to p_1 and p_i , resulting (for each k) in two curves $\Delta'_{i,k}$ and $\Delta''_{i,k}$ connecting p_1 to p_i , where $\Delta'_{i,k}$ lies outside C (these are the curves we use for G_1) and $\Delta''_{i,k}$ inside C (these are the curves we use for G_2). Each such curve has at most $2|E(T)| - 1$ bends. As in the proof of Pach and Wenger, we can create edges $p_i p_j \in E(G_1)$ by concatenating $\Delta'_{i,k}$ with $\Delta'_{j,k}$. Since we chose m_1 such approximations, we can do this for each edge in G_1 . There are two problems remaining: edges $p_i p_j$ now all pass through p_1 and they could potentially cross (rather than just touch) there. Pach and Wenger show that any two edges touch, so the drawing can be modified close to p_1 so as to separate all edges $p_i p_j$ from each other. This introduces at most one more bend per edge, so that the resulting edges have $2(2|E(T)| - 1) + 1 = 4|E(T)| - 1$ bends. Finally, note that each edge $p_i p_j$ comes close to each leaf of T (including p_1) at most twice, once for $\Delta'_{i,k}$ and once for $\Delta'_{j,k}$.

We are finally ready to complete the proof of Lemma 3. We show how to apply Lemma 6 in case G is not Hamiltonian, and not all its vertices are assigned fixed locations.

Proof of Lemma 3: By Lemma 4, we can construct a graph G' with a Hamiltonian cycle C by subdividing each edge of G at most twice, and by adding some edges, where G' has a planar embedding extending the embedding of G .

Next we deal with the issue that not all vertices lie in U , the set of vertices with fixed locations. Traverse C : whenever we encounter an edge of C with at least one endpoint not in U , contract that edge. This yields a new Hamiltonian graph G'' with $V(G'') = U$ and a planar embedding induced by the planar embedding of G' . Use Lemma 6 to construct a planar poly-line drawing of G'' at the fixed vertex locations, and ε -close to T , so that each edge of G'' has at most $4|E(T)| - 1$ bends. Each vertex $u \in U$ of G'' corresponds to a set of vertices $V_u \subseteq V(G')$ which was contracted to u , so the subgraph G'_u of G' induced by V_u is connected. Since we embedded G'' with the induced planar embedding of G' , we can now do some surgery to turn u back into G'_u .

The idea is to remove a small disc around vertex u in the drawing of G'' , and to draw G'_u inside this disc, connected to the appropriate edges leaving the disc.

This will involve introducing new vertices where edges cross into the disc. The same idea was used in [17, Theorem 2].

To this end, we define a graph G_u^+ , which consists of G'_u , a cycle C_u containing G'_u in its interior, and some further edges. Each vertex of C_u corresponds to an edge of G' “incident to” G'_u , i.e., with an end-vertex in V_u and an end-vertex not in V_u . Vertices appear in C_u in the same order as the corresponding edges incident to G'_u leave G'_u (this order also corresponds to the cyclic order of the edges incident to u in G''); each vertex of C_u corresponding to an edge e of G' is connected to the end-vertex of e in V_u . Finally, G_u^+ contains further edges that triangulate its internal faces.

Consider a small disk δ around u . We erase the part of the drawing of G'' inside δ . We construct a straight-line convex drawing of G_u^+ in which each vertex of C_u is mapped to the point in which the corresponding edge crosses the boundary of δ . This drawing always exists (and can be constructed efficiently), since G_u^+ is 2-connected and internally-triangulated. Removing the edges that triangulate the internal faces of G_u^+ completes the reintroduction of G'_u .

Overall, we added one bend to an edge with exactly one endpoint in V_u . Since an edge can have endpoints in at most two V_u , this process adds at most two bends per edge, so every edge has at most $4|E(T)| + 1$ bends. Since each edge of G was subdivided at most twice to obtain G' , each edge of G has at most $3(4|E(T)| + 1) = 12|E(T)| + 3 < 12|V(T)|$ bends. Each edge of G' comes close to each leaf of T at most twice, so each edge of G comes close to each vertex of U at most six times. This concludes the proof of Lemma 3. \square

2.3 Proof of Theorem 1

As we mentioned earlier, it is sufficient to prove the result for each face of \mathcal{H} , so fix such a face F . Let W_i , with $1 \leq i \leq b$, be the facial walks of F . We distinguish between facial walks consisting of isolated vertices, indexed by $I := \{i : |W_i| = 1\}$, and facial walks consisting of more than one vertex, with indices in $N := \{1, \dots, b\} - I$. Construct an inner ε -approximation P_ε of $\bigcup_{i \in N} W_i$ —that is, F without the isolated vertices—using Lemma 1, and let F' be the face bounded by P_ε and isolated vertices W_i , $i \in I$. For $i \in N$ let F'_i be the polygon in F' that approximates W_i . Then $|F'_i| \leq \max\{3, |W_i|\} \leq |W_i| + 1$ by Lemma 2 and the fact that W_i has size at least 2. Thus we have that $|F'| \leq \sum_{i \in N} |W_i| + |N| + |I|$.

We can triangulate F' using at most $|F'| + 2|N| + |I| - 4$ triangles, applying the following lemma with $n = |F'|$, $h_1 = |I|$, and $h_2 = |N| - 1$.

Lemma 7 (Based on O’Rourke [24, Lemma 5.2]). *Given an n -vertex polygonal region with h_1 point-holes and h_2 non-point-holes, this region can be triangulated by adding chords in time $O(n \log n)$. The resulting triangulation has $n + h_1 + 2h_2 - 2$ triangles.*

Proof. The time bound can be derived from the algorithm of O’Rourke [24, Lemma 5.1]. Consider the total sum of all angles in triangles of the triangulation. Suppose there are n_0 vertices on the outer face, $n_1 = h_1$ isolated vertices, and n_2




Fig. 3. A face F with outer facial walk W_1 and inner facial walk W_2 . (a) The 5 edges of $G - H$. (b) The inner approximation F' (heavy blue lines), a triangulation of it (fine lines), and the dual spanning tree (dashed red) with extra vertices v_1 and v_2 close to W_1 and W_2 , respectively.

vertices on non-point-holes (of which there are h_2). Then the total angle sum is $[(n_0 - 2) + 2n_1 + (n_2 + 2h_2)]\pi$ which equals $t\pi$, where t is the number of triangles. We conclude that $t = n + h_1 + 2h_2 - 2$.

We use a result of Bern and Gilbert [3] to construct a straight-line drawing of the dual of the triangulation. Bern and Gilbert place a vertex at the *incenter* of each triangle (where the angle bisectors of the triangle meet) and prove that the straight-line edge joining two vertices in adjacent triangles lies within the union of the two triangles. Now take a spanning tree T of the dual. By Lemma 7, T has $|F'| + 2|N| + |I| - 4$ vertices. For each facial walk W_i , $i \in N$, we augment T with a new leaf v_i close to W_i and inside F' ; for each facial walk W_i , $i \in I$, we add the isolated vertex of W_i to T as a new leaf v_i . This adds $|N| + |I|$ vertices to T , so the number of vertices of T is now $n_T = |F'| + 3|N| + 2|I| - 4$.

Let G_F be the embedded multigraph obtained by restricting G to vertices and edges lying inside or on the boundary of F and by contracting each facial walk W_i of F to a single vertex v_i . We can now use Lemma 3 to embed G_F along T so that vertices v_i are drawn at their fixed locations. Each edge of G_F has at most $12n_T$ bends.

We now want to connect edges in G_F to the boundary components they belong to. For facial walks W_i , $i \in I$, there is nothing to do, since we chose v_i to be the isolated vertex which is the boundary component W_i . So we may assume that we are dealing with boundary components consisting of more than one vertex. We will use the buffer between F' and F to do this. In fact, we need to split the buffer zone into two, so we apply Lemma 1 a second time to obtain an inner $\varepsilon/2$ -approximation F'' of F , so that $F' \subseteq F'' \subseteq F$. See Figure 4. Let F''_i be the polygon that approximates W_i in F'' . Note that $|F''_i| = |F'_i| \leq |W_i| + 1$. Now for each walk W_i we extend the edges ending at v_i to their endpoint on W_i . Since we maintained the cyclic order of G_F -edges at v_i , we can simply route

these edges around W_i using approximations to W_i via Lemma 1, and we can do so in $F_i - F_i''$. This adds two bends to the edge near v_i , plus at most one bend for each vertex of F_i'' except the one corresponding to the final destination vertex on W_i . In total we add at most $2 + |F_i''| - 1 \leq |W_i| + 2$ bends. There is one difficulty: there are edges of G_F that pass by v_i , separating it from the segment of F' close to v_i (which is our gate to W_i). To remedy this difficulty, we first route all of these edges around the whole obstacle W_i in the $F'' - F'$ part of the buffer, which adds $|F_i'| + 3 \leq |W_i| + 4$ bends to an edge every time it passes v_i (see Figure 4(b), note that the edge starts with one bend close to the vertex).

Now we are free to route the G_F -edges incident to v_i to their endpoints along W_i . Since an edge can pass by and/or terminate at a vertex at most six times, the number of additional bends in each edge caused by going around W_i is at most $6(|W_i| + 4) = 6|W_i| + 24$; totalling this number over all boundary components of F yields a bound of at most $6 \sum_{i \in N} |W_i| + 24|N|$ bends along the whole edge (we can ignore W_i with $i \in I$, since we do not reroute around those components). Since each G_F -edge started with $12n_T$ bends, each G_F -edge now has at most $12n_T + 6 \sum_{i \in N} |W_i| + 24|N|$ bends.

In order to derive a bound in terms of $n_H = |V(H)|$, we use:

- (1) $n_T = |F'| + 3|N| + 2|I| - 4$ (as discussed in the first part of this subsection),
- (2) $|F'| \leq \sum_{i \in N} |W_i| + |N| + |I|$ (as discussed in the first part of this subsection),
- (3) $\sum_{i \in N} |W_i| \leq 2n_H$ (which can be easily proved by induction on $|N|$, primarily, and on the number of 2-connected components of W_i , if $|N| = 1$), and
- (4) $2|N| + |I| \leq n_H$ (since each facial walk W_i with $i \in N$ consists of more than one vertex).

From (1) and (2) we get that $n_T \leq \sum_{i \in N} |W_i| + 4|N| + 3|I|$. Thus the number of bends in each G_F -edge is at most

$$\begin{aligned}
12n_T + 6 \sum_{i \in N} |W_i| + 24|N| &\leq 12 \left(\sum_{i \in N} |W_i| + 4|N| + 3|I| \right) + 6 \sum_{i \in N} |W_i| + 24|N| \\
&\leq 18 \sum_{i \in N} |W_i| + 72|N| + 36|I| \\
&\leq 18 \left(\sum_{i \in N} |W_i| \right) + 36(2|N| + |I|).
\end{aligned}$$

From (3) and (4), we conclude that each G_F -edge has at most $36n_H + 36n_H = 72n_H$ bends.

Most of the steps in the construction can be performed in linear time. Building the triangulation takes time $O(n_H \log n_H)$. The overall running time is thus bounded by the size of the resulting drawing which contains a linear number of edges each with a linear number of bends, yielding the quadratic running time.

Remark 1. The algorithm we presented in this section provides a bound better than $72n_H$ bends per edge if the subgraph H of G for which a straight-line drawing \mathcal{H} is given as part of the input is *induced*. If that is the case, then the embedded multigraph G_F defined in this section contains no self-loops; consequently, a Hamiltonian graph G'_F can be constructed in linear time by




Fig. 4. A close-up of the situation near inner facial walk W_2 . (a) After drawing G_F around the tree T (heavy dashed line), edges $1, \dots, 5$ are incident to v_2 in the correct cyclic order, but two other edges e and f pass by between v_2 and F' . (b) We add a second approximation F'' and route the edges e and f (in dashed red) around W_2 in the buffer zone between F'' and F' . (c) We route the edges incident to W_2 in the buffer zone between F and F'' .

adding vertices and edges and by subdividing edges of G_F so that each edge is subdivided by at most one new vertex (while in the general case we use two subdivision vertices per edge, see Lemma 4). This immediately allows us to improve the bounds in Lemma 3 on the number of bends per edge to $8n_T$ and on the number of times each edge comes close to each vertex u to at most four. The same analysis as above and the improved bounds of Lemma 3 allow us to upper bound the number of bends per edge in Theorem 1 by $48n_H$.

Remark 2. An improvement upon the $72n_H$ bound of Theorem 1 can be obtained by modifying the placement of v_i , for each $i \in N$, and the route of the edges that go around W_i . This modification makes the algorithm slightly more involved, so we preferred to omit it from the proof and to sketch it here. The main idea is that vertex v_i can be inserted not just at any point inside F' , but rather at a convex corner of F'_i that approximates an occurrence σ of a vertex of W_i . Then each edge that goes around v_i and has to be “wrapped around” W_i can save three bends (each time it passes by v_i) with respect to the route described in Figure 4(b). To achieve this, we bend the edge at its intersection points with F'_i and then connect it directly to the suitable approximations of the vertices next to σ along W_i . This route introduces $|F'_i| = |W_i| + 1$ new bends each time an edge passes by v_i . A similar argument can be used for the edges that terminate at some vertex of W_i . This results in each G_F -edge having at most $12n_T + 6 \sum_{i \in N} |W_i| + 6|N|$ bends. Then the same calculations described above lead to a bound of $63n_H$ bends per edge.




Fig. 5. A drawing Γ of G that greedily extends \mathcal{H} with respect to σ . Drawing \mathcal{H} consists of the black circles. The first edges $n - N - 1$ edges in σ are (black) straight-line segments. The last N edges (u_i, v_i) are (colored) polygonal lines whose bends have been made smooth to improve the readability. Only four of the latter edges are shown.

3 Extending Partial Drawings Greedily

Let G be a plane graph with a spanning subgraph H for which we have fixed a straight-line planar drawing \mathcal{H} . For a given ordering $\sigma = [e_1, \dots, e_m]$ of the edges in $G \setminus H$ we say that a drawing Γ of G *greedily extends \mathcal{H} with respect to σ* if it is obtained by drawing edges e_1, \dots, e_m in this order, so that e_i is drawn as a polygonal curve that respects the embedding of G and with the minimum number of bends, for $i = 1, \dots, m$.

Suppose σ orders the edges of $G \setminus H$ so that the edges between distinct connected components of H precede edges between vertices in the same connected component of H . For such orderings Fowler *et al.* claimed in [14] that there exists a drawing Γ of G greedily extending \mathcal{H} with respect to σ in which each edge has $O(|V(G)|)$ bends. However, in the following we confirm a claim of Schaefer [28] stating that greedy extensions do not, in general, lead to drawings with a polynomial number of bends.

Theorem 3. *For every n there exists an n -vertex plane graph G , a planar drawing \mathcal{H} of $H = (V(G), \emptyset)$, the empty spanning subgraph of G , and an order σ of the edges in G so that any drawing of G that greedily extends \mathcal{H} with respect to σ has edges with $2^{\Omega(n)}$ bends.*

Proof. We adapt an example by Kratochvíl and Matoušek [23]. Refer to Figure 5. Let $N = \lfloor \frac{n}{3} \rfloor - 6$, for any integer n . Graph H consists of n isolated vertices, name them $u_1, \dots, u_N, v_1, \dots, v_N, w_1, \dots, w_N, a, b, c, d, e, r_1, \dots, r_{n-3N-5}$. The first $n - N - 1$ edges in σ are (u_i, w_i) for $i = 1, \dots, N$, (w_i, w_{i+1}) for $i = 1, \dots, N - 1$, (r_i, r_{i+1}) for $i = 1, \dots, n - 3N - 6$, (c, w_1) , (b, c) , (c, e) , (e, d) , (a, d) , and (a, r_{n-3N-5}) . All these edges are straight-line segments in any drawing Γ of G that greedily extends \mathcal{H} with respect to σ . The last N edges in σ are $(u_1, v_1), \dots, (u_N, v_N)$ in this order.

Consider any drawing Γ of G that greedily extends \mathcal{H} with respect to σ . We claim that edge (u_i, v_i) has 2^{i-1} bends in Γ . In fact, it suffices to prove

that (u_i, v_i) has 2^{i-1} intersections with the straight-line segment ab in Γ . Indeed, (u_1, v_1) has exactly one intersection with ab in Γ . Inductively assume that (u_i, v_i) has 2^{i-1} intersections with ab in Γ ; we prove that (u_{i+1}, v_{i+1}) has 2^i intersections with ab in Γ . This proof is accomplished by following Kratochvíl and Matoušek [23] almost *verbatim*. Since (u_{i+1}, v_{i+1}) does not cross (u_i, v_i) , it has a bend b_{i+1} around v_i , i.e., inside the square defined by $u_{i-2}, w_{i-2}, w_{i-1}$, and u_{i-1} . Thus the polygonal curve representing (u_{i+1}, v_{i+1}) in Γ consists of two parts—one from u_{i+1} to b_{i+1} , the other from b_{i+1} to v_{i+1} . Both of these parts may be used as an edge joining u_i and v_i , after contracting u_{i+1} and v_{i+1} into u_i , and b_{i+1} into v_i . Hence, by induction, each of these two parts has 2^{i-1} intersections with ab , and the whole edge (u_{i+1}, v_{i+1}) has 2^i intersections with ab .

Hence, in any drawing Γ of G that greedily extends \mathcal{H} with respect to σ , one edge has $2^{N-1} = 2^{\lfloor \frac{n}{3} \rfloor - 7} \in 2^{\Omega(n)}$ bends, which concludes the proof.

We remark that the graph G in the proof of Theorem 3 is a tree, so every edge of G connects vertices in distinct connected components of H .

4 Simultaneous Planarity

Before turning to our algorithm to draw simultaneously planar graphs, we justify our claim that minimizing the number of crossings in a simultaneous planar drawing is **NP**-hard. This result follows from Cabello and Mohar’s proof of **NP**-hardness for the *anchored planarity* problem [6, Theorem 2.1], but a more direct proof of a slightly stronger result is possible by reduction from the **NP**-complete crossing number problem.

Theorem 4. *Minimizing the number of crossings in a simultaneous planar drawing of two graphs is **NP**-complete, even if one graph is the disjoint union of paths of length at most two and the other graph is a matching.*

The result is sharp in the sense that if both G_1 and G_2 are matchings, the problem is easy, since the union of two matchings is always planar.

Proof. We use the fact that the (standard) crossing number problem is **NP**-hard for cubic graphs [18]. Let G be a cubic graph with m edges. Subdivide each edge $2m$ or $2m + 1$ times (we will shortly see which). At each of the original vertices of G choose two of the incident edges, and make them part of G_1 ; the third edge at each vertex is added to G_2 . Now add the remaining edges to G_1 and G_2 so that along each path between original vertices G_1 and G_2 edges alternate. If such a path ends with two G_1 -edges or two G_2 -edges, we need to subdivide it $2m$ times to make this possible; if it ends with one G_1 -edge and one G_2 -edge, we subdivide it $2m + 1$ times. By this construction, G_1 is a disjoint union of paths of length at most two, and G_2 is a matching. Finally, the number of crossings in a simultaneous planar drawing of G_1 and G_2 is an upper bound on the crossing number of G , and, since we subdivided each edge of G sufficiently often, the two numbers are equal: starting with a crossing-minimal drawing of G , we can realize each crossing by aligning a G_1 -edge with a G_2 -edge.

We now turn to the proof the Theorem 2.

Proof of Theorem 2: We first note that it is easy to go from (ii) to (i): Suppose we have constructed, in time $O(n^2)$ a simultaneous planar drawing Γ so that private edges of G_1 and G_2 intersect at most 24 times, all edges of G_1 are straight, and all private edges of G_2 have at most $72|V(G_1)|$ bends. We add dummy vertices at the locations of the $O(n^2)$ crossings points in Γ , and then construct a straight-line drawing of the resulting planar graph on a small grid. The number of bends per edge in the new drawing is at most $72n$, since each edge in Γ intersects fewer than $3n$ edges, and each one of them at most 24 times.

We are left with the proof of (ii). That is, we have to construct in time $O(n^2)$ a simultaneous planar drawing of G in which private edges of G_1 and G_2 intersect at most 24 times, all edges of G_1 are straight, and every private edge of G_2 has at most $72|V(G_1)|$ bends.

Start with an arbitrary straight-line planar drawing Γ_1 of G_1 . We now construct a drawing Γ_2 of G_2 using an approach similar to the proof of Theorem 1. Drawing Γ_1 induces a straight-line planar drawing Γ of G . Thus, in order to determine Γ_2 , it remains to describe how to draw the private edges of G_2 . We will accomplish this independently for each face F of G .

We construct a triangulation Σ of F by using all the vertices and edges of G_1 that lie inside F , as well as some extra edges we will specify shortly. Next, we execute the same algorithm we used in the proof of Theorem 2. Namely, we construct a straight-line drawing of the dual D of Σ and we take a spanning tree T of D . For each facial walk W_i of F , we augment T with a leaf v_i close to W_i and inside F' , if $|W_i| > 1$, and coinciding with W_i , if $|W_i| = 1$; here, F' is an inner ε -approximation of F constructed as earlier. Let G_2^F be the embedded multigraph obtained by restricting G_2 to the vertices and edges inside or on the boundary of F , and by contracting each facial walk W_i of F to a single vertex v_i . We use Lemma 3 to construct a planar poly-line drawing of G_2^F that realizes the given embedding, that is ε -close to T , and in which vertices v_i maintain their fixed locations. Finally, for boundary components with $|W_i| > 1$, we reconnect edges in G_2^F to the boundary components they belong to. In order to do this, we first “wrap” the edges of G_2^F passing by a vertex v_i around W_i , and we then extend the edges of G_2^F incident to v_i to their endpoint on W_i , by routing them around W_i .

By construction every edge of G_1 is straight. By Theorem 1 every private edge of G_2 has at most $72|V(G_1)|$ bends. Also, the algorithmic steps are the same as for the proof of Theorem 1, hence the algorithm runs in $O(n^2)$ time. It remains to prove that any private edge of G_1 and any private edge of G_2 intersect at most 24 times.

Consider any private edge e of G_2 and any private edge e' of G_1 . Recall that e' is an edge of Σ . Denote by W_i and W_j the facial walks that the end-vertices of e' belong to. Edge e can only intersect edge e' in the following two situations: when passing by v_i or v_j and when passing by the point p_T in which the edge of D dual to e' crosses e' . We prove that each of these two types of intersections happens at most 12 times.

For the first type of intersections, Lemma 3 implies that edge e passes by each of v_i or v_j at most 6 times, hence at most 12 times in total.

For the second type of intersections, Lemma 4 implies that edge e is subdivided into at most three edges e_1 , e_2 , and e_3 in order to turn G_2^F into a Hamiltonian graph. For each $j = 1, 2, 3$, e_j either belongs to the Hamiltonian cycle of the subdivided G_2^F or not. In the former case, e_j is drawn as part of an $i\varepsilon/n$ -approximation Θ_i of T , as in the proof of Lemma 5, hence it crosses e' at most twice. In the latter case, e_j is composed of two parts, denoted by $\Delta'_{p,k}$ and $\Delta'_{q,k}$, or by $\Delta''_{p,k}$ and $\Delta''_{q,k}$ in the proof of Lemma 6. Each of $\Delta'_{p,k}$, $\Delta'_{q,k}$, $\Delta''_{p,k}$ and $\Delta''_{q,k}$ is part of a $k\varepsilon/(nm_1)$ -approximation of Θ'_i , which is part of Θ_i . Hence, each of $\Delta'_{p,k}$, $\Delta'_{q,k}$, $\Delta''_{p,k}$ and $\Delta''_{q,k}$ crosses e' at most twice; thus e_j crosses e' at most four times, and e crosses e' close to p_T at most 12 times. \square

5 Open Questions

We conclude with three open questions. We proved that if a graph has a planar drawing extending a straight-line planar drawing of a subgraph then there is such a drawing with at most $72n$ bends per edge. This is asymptotically tight, but can the constant 72 be reduced? As sketched at the end of Section 2, a variation of our algorithm decreases this constant to 63, however new ideas seem to be needed in order to push the bound further down.

Our second result was that any two simultaneously planar graphs have a simultaneous planar drawing with at most 24 crossings per pair of edges and a linear number of bends per edge with a drawing on a polynomial-sized grid. The only lower bound on the number of crossings between two edges in a simultaneous planar drawing is 2 (see [9] or the figure in the margin for the entry “simultaneous crossing number” in [27]). There is a large gap between 2 and 24. Can two edges be forced to cross more than twice in a simultaneous planar drawing? For the third open question, we note that Grilli et al. [16] showed that two simultaneously planar graphs have a drawing with at most 9 bends per edge, though with a larger constant for the number of crossings and not on a grid. Is it possible to achieve the best of both results: 9 bends per edge, 24 crossings per pair of edges, and a nice grid?

References

1. P. Angelini, G. Di Battista, F. Frati, V. Jelínek, J. Kratochvíl, M. Patrignani, and I. Rutter. Testing planarity of partially embedded graphs. In *Proc. Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010*, pages 202–221. SIAM, 2010. URL: <http://dl.acm.org/citation.cfm?id=1873601>. 1, 3, 4
2. M. Badent, E. Di Giacomo, and G. Liotta. Drawing colored graphs on colored points. *Theor. Comput. Sci.*, 408(2-3):129–142, 2008. 4
3. M. Bern and J. R. Gilbert. Drawing the planar dual. *Inform. Process. Lett.*, 43(1):7–13, 1992. URL: [http://dx.doi.org/10.1016/0020-0190\(92\)90022-N](http://dx.doi.org/10.1016/0020-0190(92)90022-N), doi:10.1016/0020-0190(92)90022-N. 13

4. T. C. Biedl and P. Floderus. Drawing planar graphs on points inside a polygon. In *Mathematical Foundations of Computer Science (MFCS 2012)*, volume 7464 of *Lecture Notes in Computer Science*, pages 172–183. Springer, 2012. URL: http://dx.doi.org/10.1007/978-3-642-32589-2_18, doi:10.1007/978-3-642-32589-2_18. 4
5. T. Bläsius, S. G. Kobourov, and I. Rutter. Simultaneous embeddings of planar graphs. In R. Tamassia, editor, *Handbook of Graph Drawing and Visualization*, Discrete Mathematics and Its Applications, chapter 11, pages 349–382. Chapman and Hall/CRC, 2013. 2, 3
6. S. Cabello and B. Mohar. Adding one edge to planar graphs makes crossing number and 1-planarity hard. *SIAM Journal on Computing*, 42(5):1803–1829, 2013. URL: <http://epubs.siam.org/doi/abs/10.1137/120872310>, arXiv:<http://epubs.siam.org/doi/pdf/10.1137/120872310>, doi:10.1137/120872310. 4, 17
7. T. M. Chan, F. Frati, C. Gutwenger, A. Lubiw, P. Mutzel, and M. Schaefer. Drawing partially embedded and simultaneously planar graphs. In C. Duncan and A. Symvonis, editors, *22nd International Symposium on Graph Drawing (GD '14)*, 2014. To appear. 1
8. T. M. Chan, H.-F. Hoffmann, S. Kiazzyk, and A. Lubiw. Minimum length embedding of planar graphs at fixed vertex locations. In S. Wismath and A. Wolff, editors, *Graph Drawing*, volume 8242 of *LNCS*, pages 376–387. Springer, 2013. URL: http://dx.doi.org/10.1007/978-3-319-03841-4_9. 9, 10
9. M. Chimani, M. Jünger, and M. Schulz. Crossing minimization meets simultaneous drawing. In *PacificVis*, pages 33–40. IEEE, 2008. 19
10. E. Di Giacomo, W. Didimo, G. Liotta, H. Meijer, and S. K. Wismath. Point-set embeddings of trees with given partial drawings. *Computational Geometry: Theory and Applications*, 42(6-7):664–676, 2009. 4
11. E. Di Giacomo and G. Liotta. Simultaneous embedding of outerplanar graphs, paths, and cycles. *Int. J. Comput. Geometry Appl.*, 17(2):139–160, 2007. 2, 5
12. C. Erten and S. G. Kobourov. Simultaneous embedding of planar graphs with few bends. *J. Graph Algorithms and Appl.*, 9(3):347–364, 2005. doi:10.1007/978-3-540-31843-9_21. 2, 5
13. A. Estrella-Balderrama, E. Gassner, M. Jünger, M. Percan, M. Schaefer, and M. Schulz. Simultaneous geometric graph embeddings. In S.-H. Hong, T. Nishizeki, and W. Quan, editors, *Graph Drawing*, volume 4875 of *LNCS*, pages 280–290. Springer, 2008. 4
14. J. J. Fowler, M. Jünger, S. G. Kobourov, and M. Schulz. Characterizations of restricted pairs of planar graphs allowing simultaneous embedding with fixed edges. *Comput. Geom.*, 44(8):385–398, 2011. 2, 16
15. E. Gassner, M. Jünger, M. Percan, M. Schaefer, and M. Schulz. Simultaneous graph embeddings with fixed edges. In F. V. Fomin, editor, *Graph-Theoretic Concepts in Computer Science*, volume 4271 of *LNCS*, pages 325–335. Springer, 2006. 2
16. L. Grilli, S.-H. Hong, J. Kratochvíl, and I. Rutter. Drawing simultaneously embedded graphs with few bends. In *Graph Drawing (GD '14)*, 2014. 3, 19
17. B. Haeupler, K. R. Jampani, and A. Lubiw. Testing simultaneous planarity when the common graph is 2-connected. *J. Graph Algorithms and Appl.*, 17(3):147–171, 2013. URL: <http://jgaa.info/accepted/2013/HaeuplerJampaniLubiw2013.17.3.pdf>, doi:10.7155/jgaa.00289. 5, 12
18. P. Hliněný. Crossing number is hard for cubic graphs. *J. Combin. Theory Ser. B*, 96(4):455–471, 2006. 17
19. V. Jelínek, J. Kratochvíl, and I. Rutter. A Kuratowski-type theorem for planarity of partially embedded graphs. *Comput. Geom.*, 46(4):466–492, 2013. 3

20. M. Jünger and M. Schulz. Intersection graphs in simultaneous embedding with fixed edges. *J. Graph Algorithms Appl.*, 13(2):205–218, 2009. 2, 5, 6
21. F. Kammer. Simultaneous embedding with two bends per edge in polynomial area. In L. Arge and R. Freivalds, editors, *10th Scandinavian Workshop on Algorithms and Theory (SWAT '06)*, pages 255–267, 2006. 2, 5
22. M. Kaufmann and R. Wiese. Embedding vertices at points: Few bends suffice for planar graphs. *J. Graph Algorithms and Appl.*, 6(1):115129, 2002. 5
23. J. Kratochvíl and J. Matoušek. String graphs requiring exponential representations. *J. Comb. Theory, Ser. B*, 53(1):1–4, 1991. 16, 17
24. J. O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, NY, 1987. 12
25. J. Pach and R. Wenger. Embedding planar graphs at fixed vertex locations. *Graphs Combin.*, 17(4):717–728, 2001. 1, 2, 4, 6, 7, 9
26. M. Patrignani. On extending a partial straight-line drawing. *Internat. J. Found. Comput. Sci.*, 17(5):1061–1069, 2006. 4
27. M. Schaefer. The graph crossing number and its variants: A survey. *The Electronic Journal of Combinatorics*, 20:1–90, 2013. Dynamic Survey, #DS21. 19
28. M. Schaefer. Toward a theory of planarity: Hanani-Tutte and planarity variants. *J. of Graph Algorithms and Appl.*, 17(4):367–440, 2013. 16